Return To Agenda Page

GALS Interconnect:
Delivering Transparent Connectivity for Multicore SoCs

Uri Cummings
VP, Product Development

FULCRUM microsystems

Fulcrum Snapshot

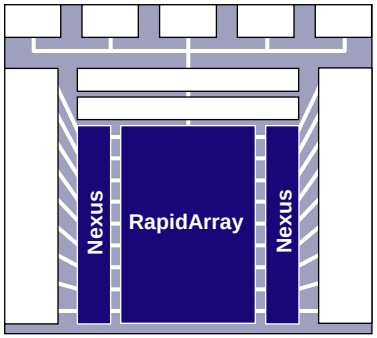

“Clockless” Semiconductor Company

Formed out of Caltech (January 2000)

Shipping low-latency 10G products today

Located in Calabasas, CA (45 people)

GraniteVentures

infinity | capital

Palomar VENTURES

NEA New Enterprise Associates

WORLDVIEW TECHNOLOGY PARTNERS

Backed by top-tier investors

Agenda

Multicore Interconnect Requirements
Limitations of Traditional Architectures
Introduction to GALS Architectures
Introduction to Nexus Interconnect
Application Examples

Multicore Interconnect Requirements

The ultimate goal: "Transparent connectivity"

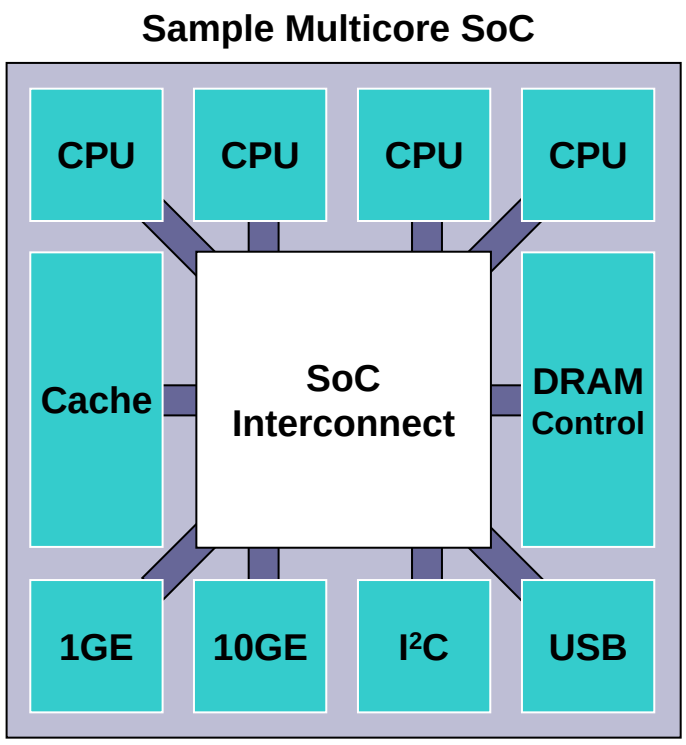

- Any-to-any connectivity
- Any number of blocks
- Massive bandwidth
- Negligible latency penalty
- Modest power and area
- Transports any type of data
- Plug-and-play design
- Fully reusable
- Process portable

Traditional Architectures Come Up Short

Fully custom, or sub-optimal performance

Split Bus Architecture

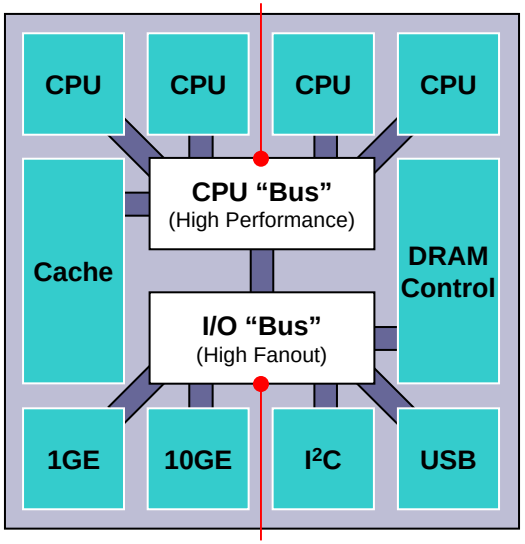

- **Specialized per design**
  - Not reusable
- **Capacity and connectivity limited to intended flow**
- **Latency varies by link**
- **Not "any-to-any"**

Shared Bus Architecture

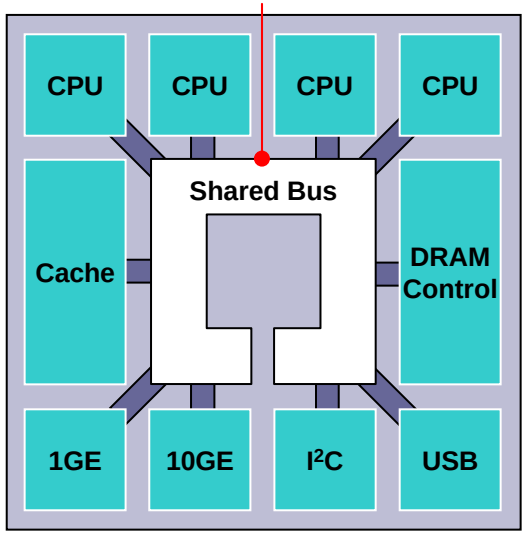

- **Under provisioned**
  - Or power hungry
- **Block placement impacts overall performance**
  - A balanced bus offers better performance

Synchronous Crossbar Architecture

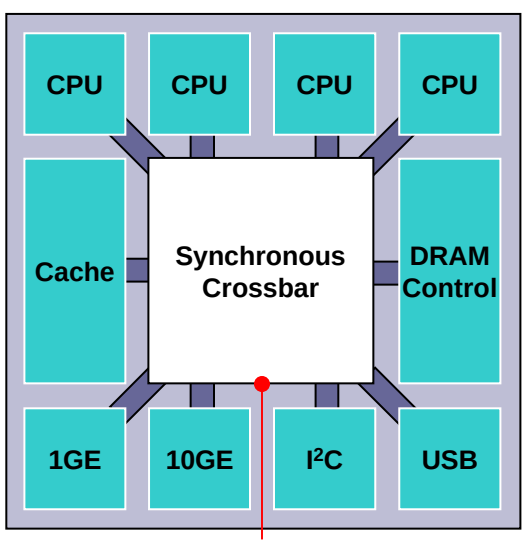

- **Power hungry**
- **Two domain hops per link**
  - Increased latency
  - Reduced performance

FULCRUM microsystems

GALS: A Novel Architectural Approach

Autonomous blocks linked together in a decoupled fashion

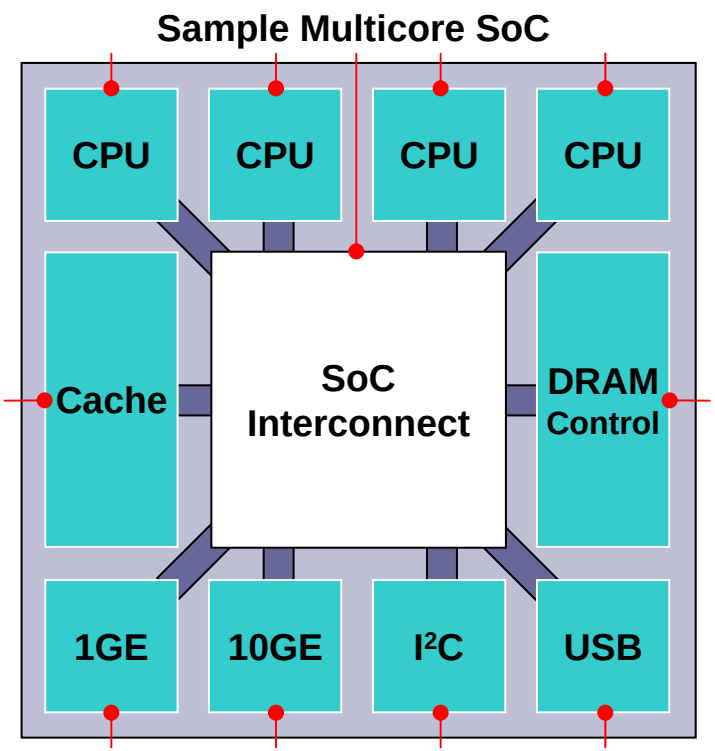

- **Simplifies large-scale design**
  - Communication exclusively through the interconnect
  - Autonomous design teams can more easily collaborate
- **Increases local performance**
  - Blocks are independently performance-optimized
  - No global timing constraints
- **Delivers reusability**
  - Blocks are self contained
  - Interconnected via generic request/grant FIFO interface
- **Improves chip performance**

GALS: Globally asynchronous, locally synchronous

GALS Interconnect Requirements

*Amended goal: "Transparent **decoupled** connectivity"*

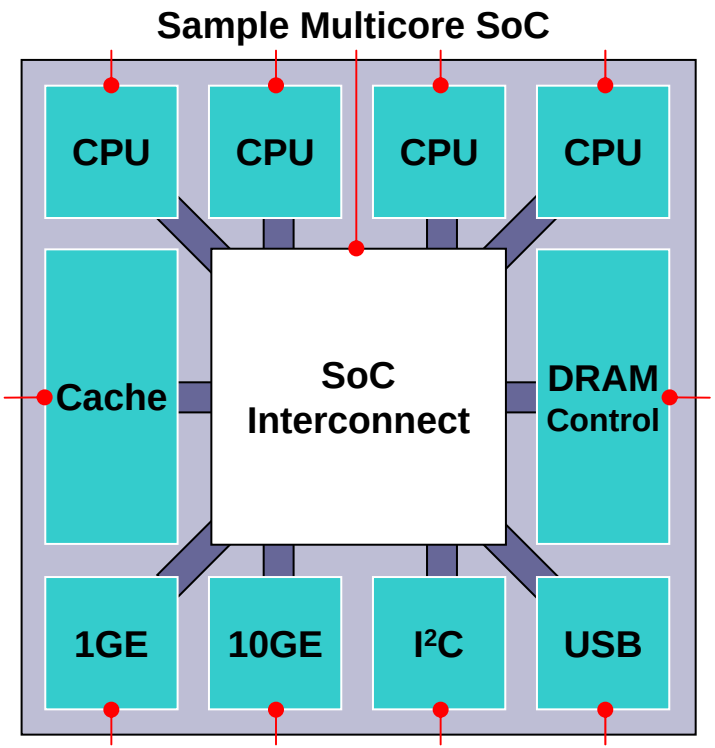

- **Block autonomy**
  - Any-to-any connectivity
  - Any number of blocks
  - Massive bandwidth
  - Negligible latency penalty
  - Modest power and area
  - Transports any data
  - Plug-and-play design
  - Fully reusable
  - Process portable

Fulcrum's Nexus Meets All Requirements

Nexus fits any performance-optimized application

Sample Multicore SoC (Based on Nexus)

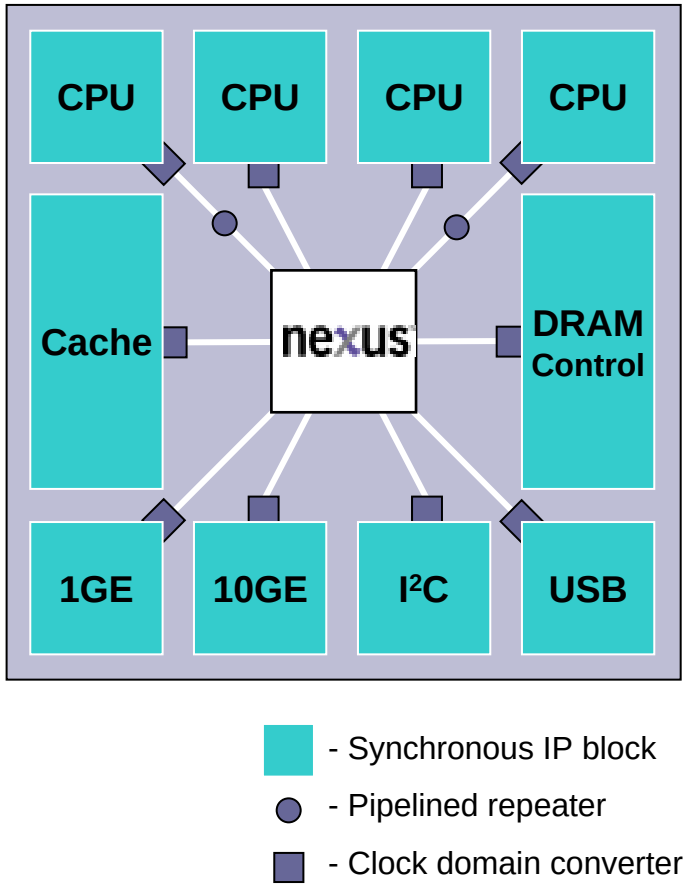

- **Massive bandwidth**
  - Up to 160Gbps/port (>2.5Tbps total)
- **Negligible latency penalty**
  - 3ns (includes arbitration)
- **Modest power and area**
  - Power scales on activity (300Gb/W)
  - Extremely compact (<2mm^2)

FULCRUM microsystems

Fulcrum's Nexus Meets All Requirements

*Nexus delivers "transparent d**ecoupled** interconnect"*

Sample Multicore SoC (Based on Nexus)

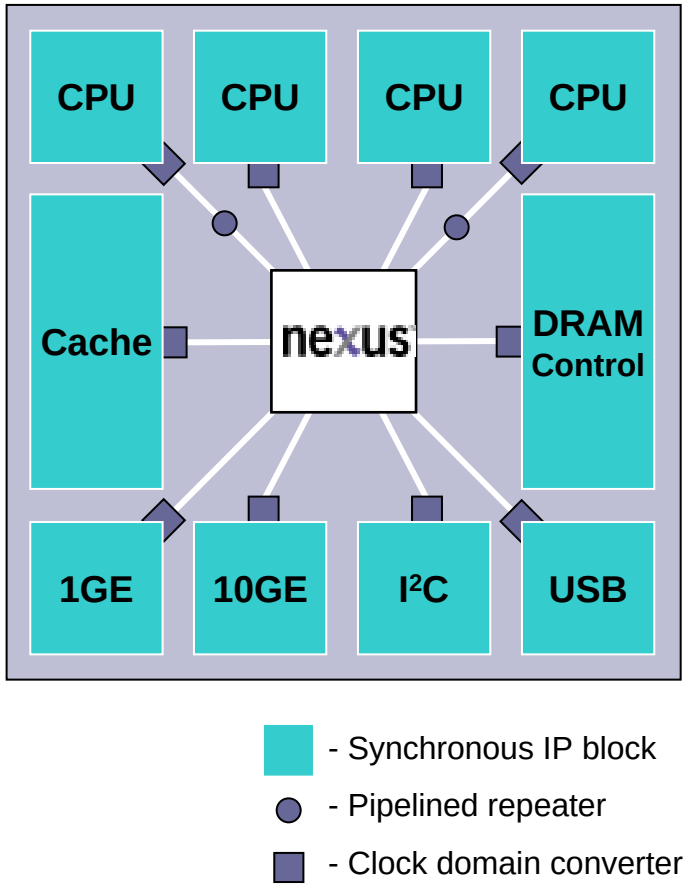

- **Block autonomy**
  - Transparently interconnects domains (any-to-any)
- **Any number of blocks**
  - Scales in ports and bit width
- **Transports any type of data**
  - Cells, packets, media streams
- **Plug-and-play design**
  - Request/grant FIFO interface
- **Fully reusable**
  - Macros placed as needed
- **Process portable**
  - 180nm, 150nm, 130nm, 90nm designs

FULCRUM microsystems

Issues With High-Speed Clocked Circuits

Worst-case performance; constant power drain

Conventional Clocked Circuits

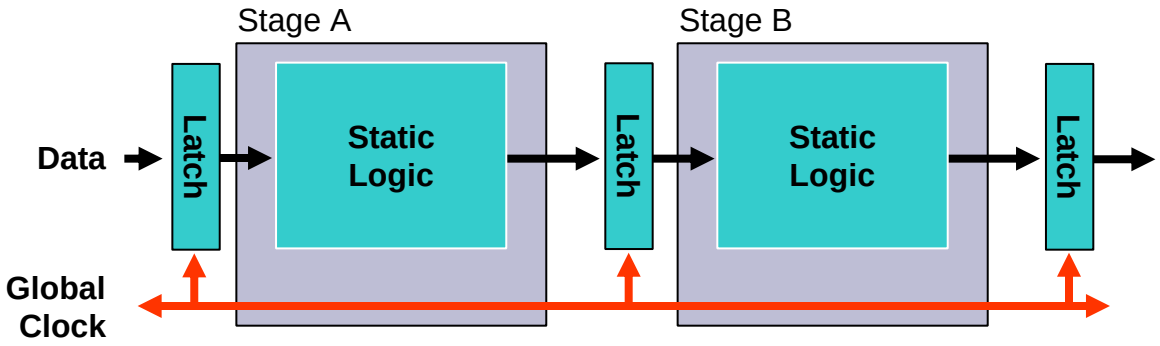

Performance
- Clock set for slowest block — penalizing overall performance

Power
- Clock structure constantly consumes power; clock gating is time consuming

Reliability
- High-speed designs are subject to races or glitches from variations in process or temp

FULCRUM microsystems

Fulcrum's Integrated Pipelining

The fastest datapath in CMOS (proven in silicon)

Fulcrum Integrated Pipelining

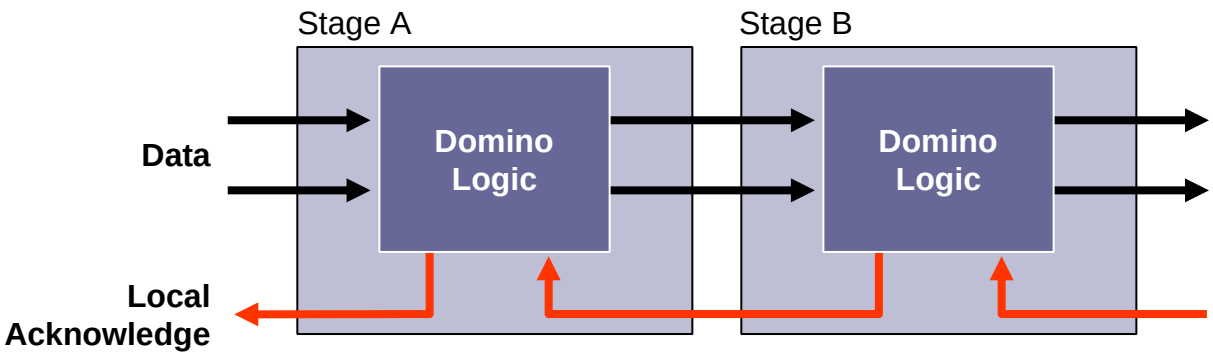

High performance
- Circuits operate with no restrictions — no clock to throttle performance

Low power
- Power consumption based on activity — no clock tree

Highly reliable
- Local handshake is robust to change — no races or glitches from variations in process and temperature

Unique GALS SoC Enabling Circuits

Two key IP blocks enable SoC performance and scalability

Nexus*

(Terabit Asynchronous Crossbar)

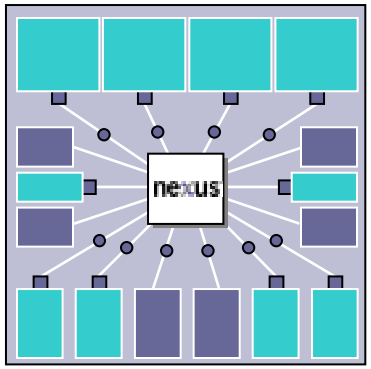

- Gigahertz performance
- Terabit capacity
- Nanosecond latency
- No power penalty

* Licensed to PMC-SIERRA for SoC interconnect

RapidArray

(Asynchronous SRAM)

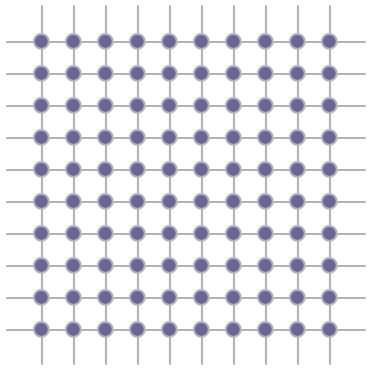

- 600MHz SRAM
- Gigahertz interconnect
- 518Gbps throughput
- Scalable for any use

Key Benefits:

- **Easily integrates independent clock domains**
- **Interconnects everything**
- **Power scales on activity**
- **2x the speed of vendor cores (same size, density, yield)**
- **Reduces power consumption (based on activity)**

FULCRUM microsystems

Nexus in a Dual-Core SoC

A real-world example

Sample Multicore SoC (Based on Nexus)

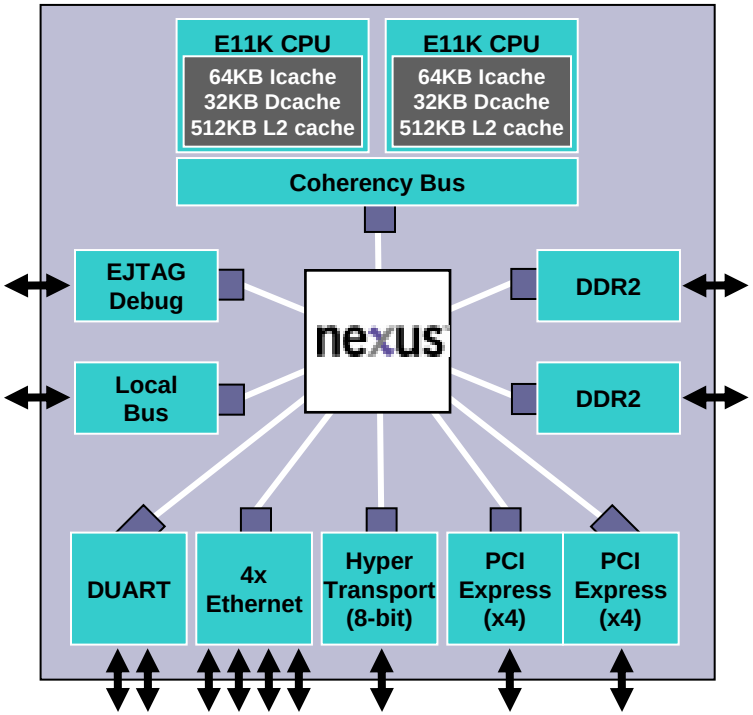

- **Dual-core MIPS processor**
  - MIPS E11K 64-bit CPUs
  - 1.8GHz
- **Uses Nexus Crossbar**
  - 90nm process technology
  - 16 ports
  - 80 bits per port
  - 1.8GHz
  - 2.3Tbps total capacity
  - <3ns cross-chip latency
  - <2mm^2 die area

FULCRUM microsystems

New-Generation Multicore System

Asynchronous crossbar provides coherent connectivity

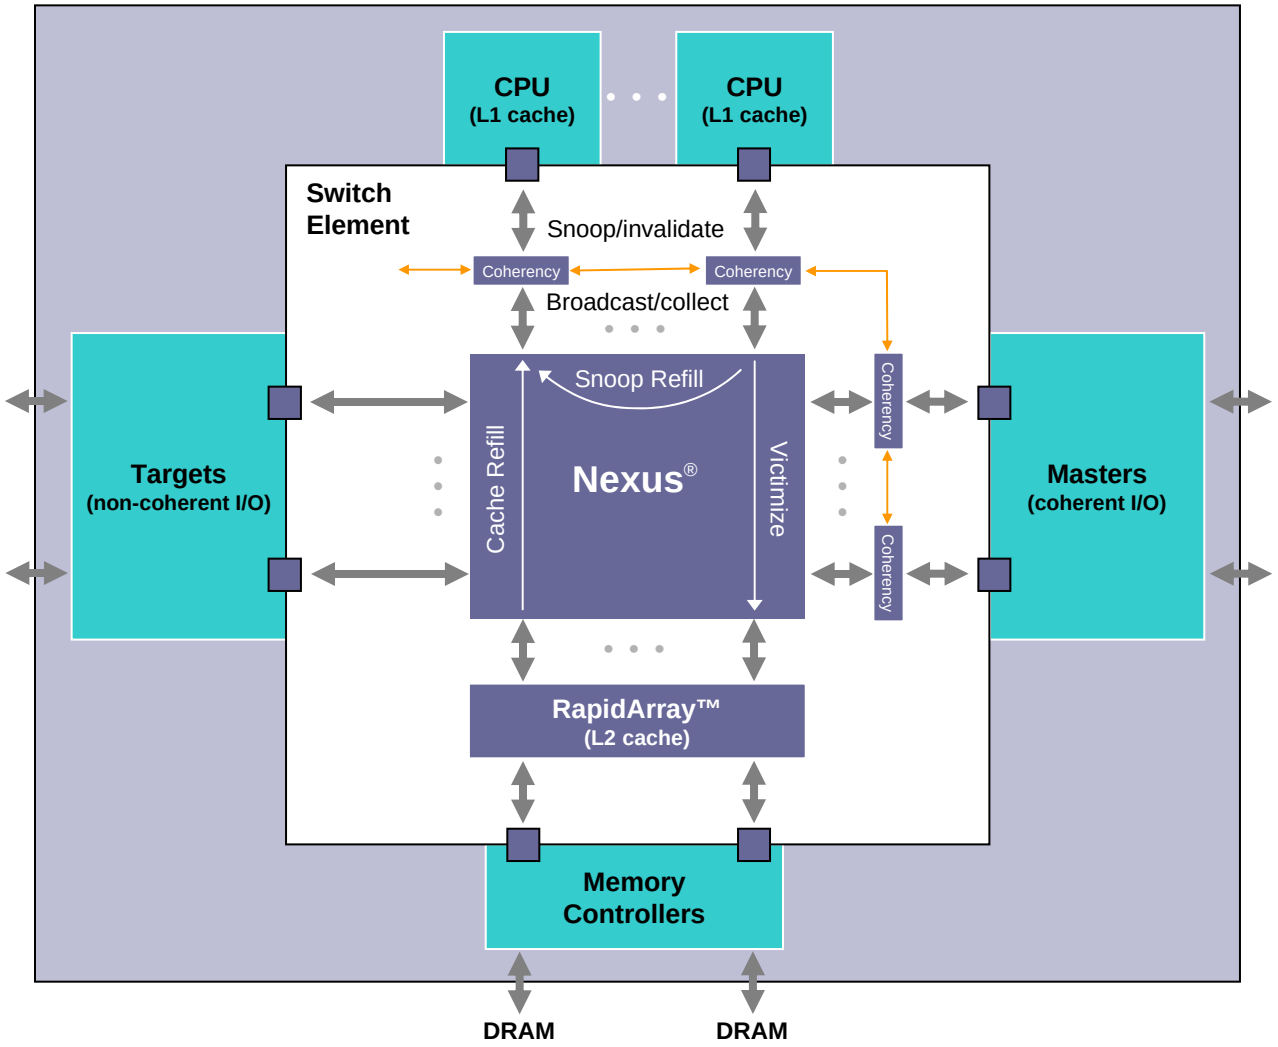

- **All system elements scale independently**
- **RapidArray can be a shared L2 cache**
- **GALS-based design eases integration**
- **Coherency can be extended to any system element**

FULCRUM microsystems

Nexus in a Commercial Chip (PivotPoint)

A real-world GALS example (13 timing domains)

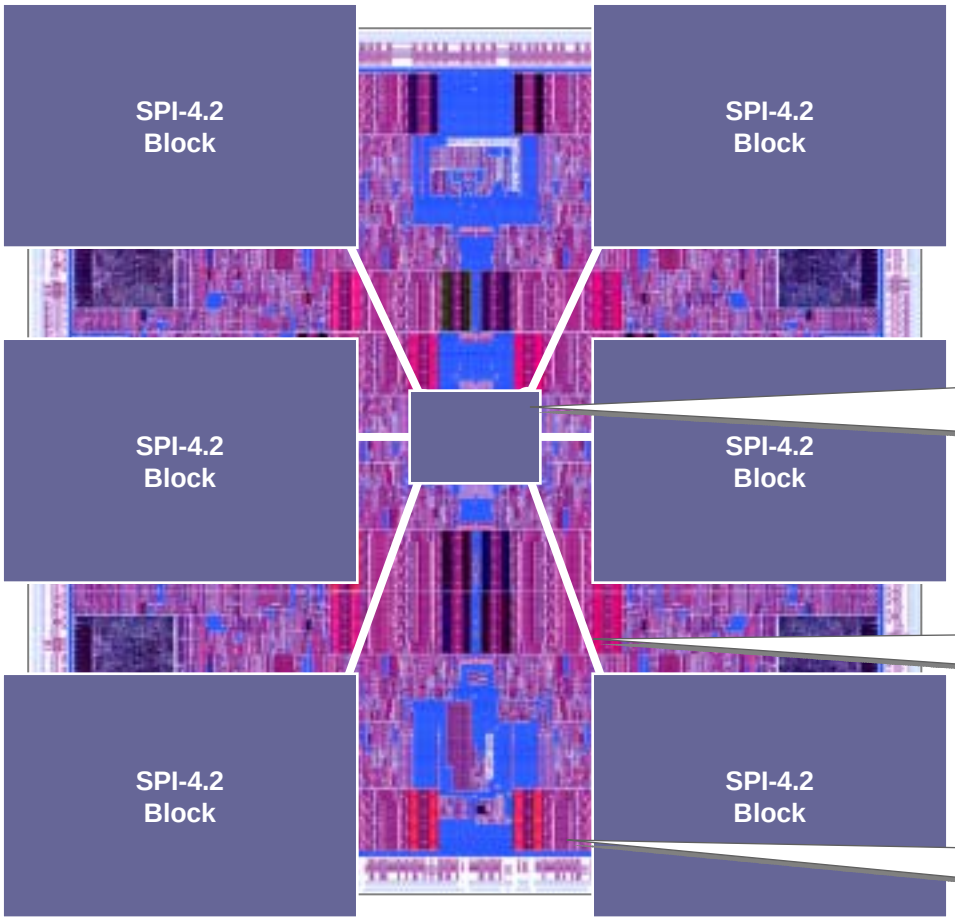

PivotPoint
- SPI-4.2 switch chip
- 14.4Gbps per port
- 13 timing domains

Nexus in PivotPoint
- 1.0GHz in TSMC 130nm
- 6 ports (36 bits wide)
- 36Gbps/port
- 384Gbps total capacity
- <3ns (port-to-port)

Async Channel
- 36Gbps/port
- Arbitrary distance
- Negligible latency

RapidArray Async SRAM
- ~200KB
- 600MHz
- Quad port (virtual)

FULCRUM microsystems

Tahoe: Massive Capacity Ethernet Switch

Shared memory architecture (dozens of timing domains)

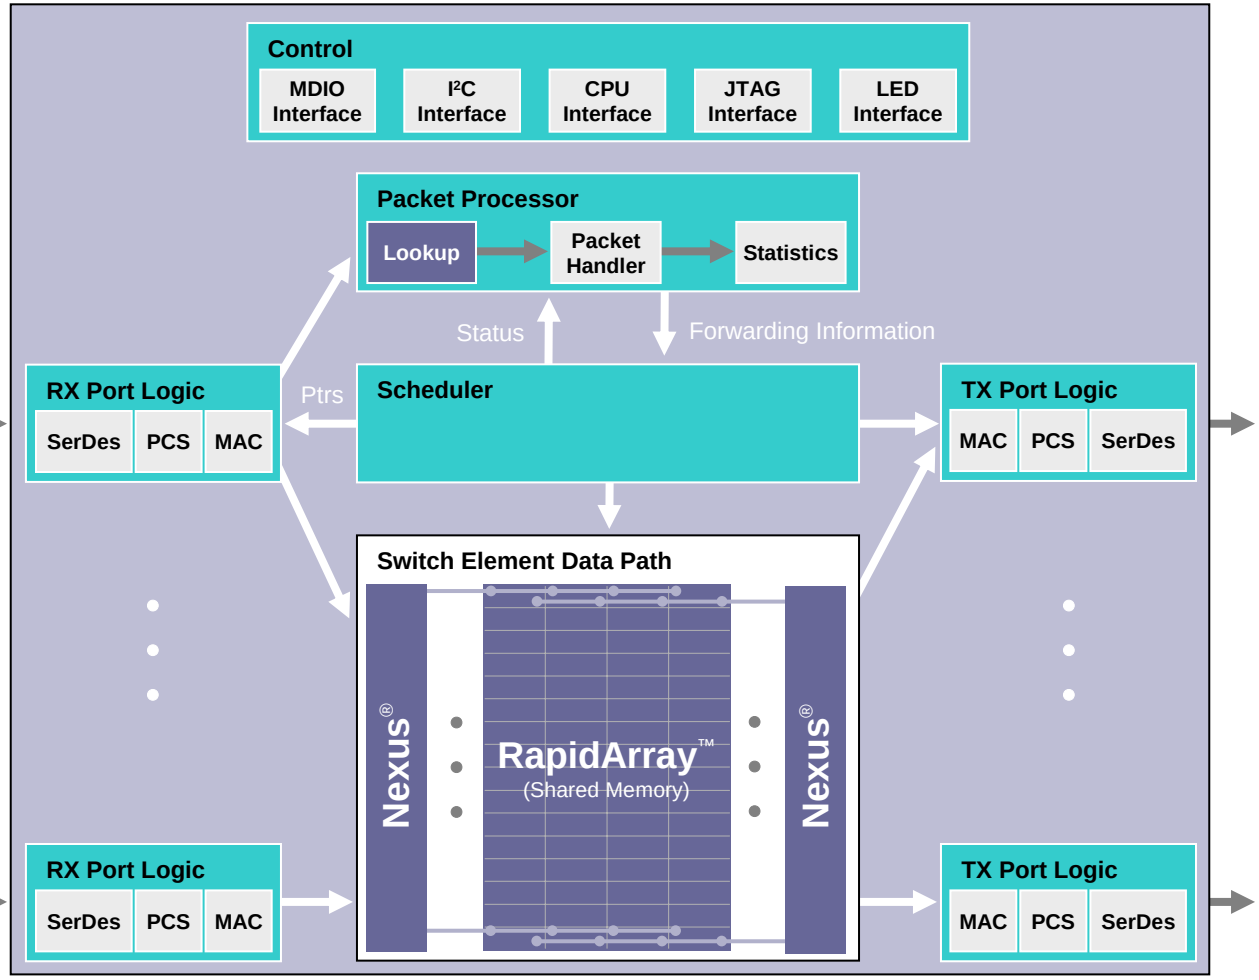

- **Local performance sized to function**
  - GHz datapath
  - 400MHz ports
  - 66MHz control
- **All ports independent**
  - Variable rates
  - Robust and resilient
- **Scalable architecture**
  - Deep roadmap with quick derivatives
- **Industry leading:**
  - Latency
  - Density
  - Packet rate
  - Power profile

FULCRUM microsystems

Thank You

Uri Cummings

VP Product Development

uri@**fulcrum**micro.com

FULCRUM microsystems

818.871.8100

www.**fulcrum**micro.com

26775 Malibu Hills Road
Suite 200
Calabasas Hills, CA 91301

"A group of engineers wants to turn the microprocessor world on its head by doing the unthinkable: tossing out the clock and letting the signals move about unencumbered. For those designers, inspired by research conducted at Caltech, **clocks are for wimps**."

Anthony Cataldo , EE Times